

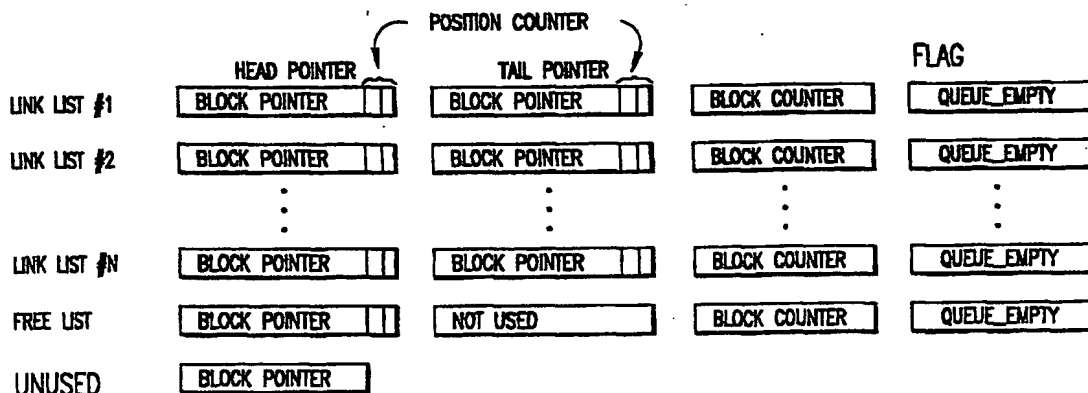


B3

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 12/00		A1	(11) International Publication Number: WO 98/36357
			(43) International Publication Date: 20 August 1998 (20.08.98)
(21) International Application Number: PCT/US98/02131		(81) Designated States: CA, CN, IL, JP, MX, NO, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 5 February 1998 (05.02.98)			
(30) Priority Data: 08/796,085 5 February 1997 (05.02.97) US		Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.	
(71) Applicant: TRANSWITCH CORPORATION [US/US]; 3 Enterprise Drive, Shelton, CT 06484 (US).			
(72) Inventors: LAU, Joseph, C.; 1F, 29 Bamboo Road III, Science-based Industry Park, Hsinchu (TW). ROY, Subhash, C.; Apartment 3A, 905 Mix Avenue, Hamden, CT 06514 (US). CALLAERTS, Dirk, L., M.; Hoevestraat 13, B-3110 Rotselaar (BE). VANDEWEERD, Ivo, Edmond, Nicole; Vuurkruisenlaan 1, B-3500 Hasselt (BE).			
(74) Agent: GORDON, David, P.; 65 Woods End Road, Stamford, CT 06905 (US).			

(54) Title: SHARED MEMORY CONTROL USING MULTIPLE LINKED LISTS WITH POINTERS, STATUS FLAGS, MEMORY BLOCK COUNTERS AND PARITY



(57) Abstract

Apparatus and methods for allocating shared memory utilizing linked lists (LLs) use a management RAM which controls the flow of data to/from a shared memory (RAM), and stores information regarding a number of LLs and a free link list (FLL) in the RAM, and a block pointer to unused RAM locations. A head pointer (HP), tail pointer (TP), block counter and empty flag (EF) are stored for each data link list. The HP and TP each include a block pointer and a position counter. The block counter contains the number of blocks used in the particular queue. An EF indicates an empty queue. The FLL includes a HP, a block counter, and an EF. Each page of RAM receiving the incoming data includes locations for storing data. The last location of the last page in a block stores a next-block pointer plus parity information, and in the last block of a queue, is set to all ones. An independent agent used in the background monitors the integrity of the LL structure.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

SHARED MEMORY CONTROL USING MULTIPLE LINKED LISTS WITH POINTERS, STATUS FLAGS, MEMORY BLOCK COUNTERS AND PARITY

This application is related to co-owned U.S. Serial No. 08/650,910, filed May 17, 1996, which is hereby incorporated by reference herein in its entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to memory management. More particularly, the present invention relates to apparatus and methods of managing a plurality of data queues stored in linked lists in a shared common memory. The invention has particular application to the use of a very large scale integrated circuit (VLSI) for the buffering of telecommunications information such as ATM data, although it is not limited thereto.

2. State of the Art

In high speed communication networks, the management of buffer resources is one mechanism of increasing network performance. One group of methods of managing buffer resources is known as sharing, where a single RAM is simultaneously utilized as a buffer by a plurality of different channels. Various sharing methods are known (see Velamuri, R. et al., "A Multi-Queue Flexible Buffer Manager Architecture", IEEE Document No. 0-7803-0917-0/93) and each has inherent advantages coupled with inherent disadvantages in terms of blocking probability, utilization, throughput, and delay. What is common to all sharing methods, however, is that a mechanism is required to direct data into appropriate locations in the RAM in a desired order so that the data can be retrieved from the RAM appropriately. One such mechanism which is well known is the use of link lists which are used to manage multiple queues sharing a common memory buffer. Typically, a link list comprises bytes of data, where each byte has at least one pointer (forward and/or backward) attached to it, thereby identifying the location of the next byte of data in the queue. The link list typically includes extensive

initialization and self-check procedures which are carried out by a microprocessor on a non-real-time basis. Thus, the use of standard prior art link list structures to manage multiplex queues sharing a common memory is not readily adaptable for VLSI implementation, and is likewise not particularly suited to the handling of very high speed telecommunications information where processing and handling are dictated by the data rate of the real-time telecommunications signal.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide an apparatus and method for control of memory allocation.

It is another object of the invention to provide a new link list structure for managing queues in a shared memory.

It is a further object of the invention to provide a single VLSI which utilizes a link list structure for managing queues of high speed real time data in a shared memory.

It is an additional object of the invention to provide a link list apparatus and method for controlling the flow of Asynchronous Transfer Mode (ATM) telecommunications data into and out of a shared buffer.

Another object of the invention is to provide an apparatus and method for VLSI control of ATM data into and out of a shared RAM by utilizing a separate RAM containing information related to the plurality of link lists in the shared RAM.

In accord with the objects of the invention a management RAM contained within a VLSI is provided for controlling the flow of data into and out of a shared memory (data RAM). The management RAM is preferably structured as an x by y bit RAM which stores information regarding y-2 data link lists in the shared RAM, a free link list in the shared RAM, and a block pointer to unused shared RAM locations. Information stored in the x bits for each

data link list includes a head pointer, a tail pointer, a block counter and an empty flag. In a preferred embodiment particularly applicable to the control of ATM data, the head and tail pointers are each composed of a block pointer and a position counter, with the position counter indicating a specific page in a block which is made up of a set of contiguous pages of memory, and the block pointer pointing to the block number. Regardless of how constituted, the head pointer contains the address of the first word of the first memory page of the link list, and the tail pointer preferably contains the address of the first word of the last memory page in the link list. The block counter contains the number of blocks used in the particular queue, and has a non-zero value if at least one page is used in the queue. The empty flag indicates whether the queue is empty such that the content of the link list should be ignored if the queue-empty flag indicates that the queue is empty.

Information stored in the management RAM for the free link list includes a head pointer, a block counter, and an empty flag, but does not need to include a tail pointer as free blocks are added to the top of the free list according to the preferred embodiment of the invention. As is discussed below in more detail, as data from different channels is directed into blocks of the data RAM, a link list is kept for each channel. As data is read out of the data RAM, blocks become available to receive new data. It is these freed blocks which are added to the free list. Block space can be assigned from the free list before or after the unused blocks (discussed below) are used.

To avoid excessive initialization requirements, an unused-block pointer is provided in the management RAM, as discussed above, and provides a pointer to the next unused block in memory. Initially all link lists, including the free list, are empty, and the unused block pointer is set to the number of blocks in the memory. As data is written to a block of shared RAM memory, the unused block pointer is decremented. When the unused block pointer equals zero, all of the cell blocks are included in the link lists (including the free link list).

According to a preferred aspect of the invention, each memory page of the shared data RAM receiving the incoming data (which RAM is managed by the management RAM) is composed of M contiguous memory addresses. Depending on the memory type, each address location can be of size B bits. The most common sizes are eight bits (byte), sixteen bits (word), thirty-two bits, and sixty-four bits. The first M-1 locations in the page are used to store data. The last (M'th) location of the last page in the block preferably is used to store the address of the first location of the next block of the queue plus an odd parity bit; i.e., the M'th location of the last page in the block stores a next block pointer plus parity information. If there are no more blocks in the queue, the M'th location in the last page is set to all ones.

According to another aspect of the invention, an independent agent is utilized in the background to monitor the integrity of the link list structure. The independent agent monitors the sum of the count of all of the link list block counters plus the unused blocks to ensure that it equals the total number of memory blocks in the common RAM. If not, an error is declared. Likewise, the independent agent checks each link list stored in the management RAM for the following error conditions: head and tail pointers are equal and the block counter is not of value one; head and tail pointers are different and the block counter is one; and, block counter equals zero. If desired, the independent agent can also monitor the block pointers stored in the M'th location of the last page of each block to determine parity errors and/or to determine errors using parity or CRC.

Using the methods and apparatus of the invention, four operations are defined for ATM cell management: cell write, cell read, queue clear, and link list monitoring. In the cell write operation, a cell is stored into a queue. More particularly, when an ATM cell is received at a port w so that it is to be stored in queue number n (which stores cells of priority v for port w), a determination is first made as to whether the queue is empty. If it is not empty, the queue status (i.e., the tail

pointer and position counter stored in management RAM) is obtained, and a determination is made as to whether a new block will be needed to be added to the queue. If a new block is not required, the cell is written to the location indicated by the tail pointer position, and the tail pointer position counter for that queue in the management RAM is updated. If this is the last page of a block, the M'th location of the page (in the shared memory) is set to all ones. If a new block is required, either because the queue was empty or because a previous cell had been written into the last page of a block, a block must be obtained. If it is a first block of a queue, initial queue parameters are stored. If it is not the first block of the link list, a block is obtained from the free list and the free list is updated; or the block is obtained from the unused blocks and the block pointer for the unused blocks is updated. Then, the cell is written to the queue, and the tail pointer, position counter, and block counter for the queue are all updated in the management RAM.

The cell read operation is utilized where a cell is to be read from a queue. In the cell read operation, the cell indicated by the head pointer and head pointer position counter for that queue is read from the queue. After reading the cell from the queue a determination is made as to whether the cell was either the last cell in a block and/or the last cell in the queue. If it is neither, then the queue status is updated (i.e., the head pointer position counter is changed), and another cell read operation is awaited. If the cell is the last cell in the block, then the queue status preferably is checked for correctness by verifying the parity of the pointer (using a parity bit), and is updated by changing the head pointer and head pointer position counter. The free list is updated by adding the freed block to the head of the free list, and the free list and link list block counters are updated. If the cell is the last cell in the queue, the procedure for the last cell in the block is followed, and the queue empty flag is set.

The queue clear operation is a microprocessor command provided for the purpose of clearing a queue. When the queue clear operation is presented, the queue status is updated by setting the queue flag, and the blocks in the queue are added to the head of the free list which is likewise updated.

The link list monitoring operation is the agent which monitors the integrity of the link list structure whenever the cell write, cell read, and queue clear operations are not running. As set forth above, the link list monitoring operation monitors the linked lists for errors by checking that the sum of the count of all of the link list block counters plus the unused blocks equals the total number of memory blocks in the common RAM, that when head and tail pointers are equal the block counter is set to one, that when head and tail pointers are different the block counter is not set to one, etc.

Additional objects and advantages of the invention will become apparent to those skilled in the art upon reference to the detailed description taken in conjunction with the provided figures.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of an apparatus incorporating the link list memory management RAM of the invention.

Figure 2 is a chart showing the structure of the memory management RAM of Figure 1.

Figure 3a is a diagram of an example of the shared data memory of the apparatus of Figure 1.

Figure 3b is a diagram of the details of a page of one of the blocks shown in Figure 3a.

Figure 3c is a diagram of an example of the information contained in the memory management RAM of Fig. 1 for managing the shared data memory example of Figure 3a.

Figures 4a - 4d are flow charts for the write, read, queue clear, and link list monitoring operations carried out by the flow controller of the apparatus of Figure 1.

Figures 5a-5d are state machine diagrams for a write, read, clear, and monitor state machine according to the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

The invention will now be described with reference to the physical layer VLSI portion of an ATM destination switch described in parent U.S. Serial No. 08/650,910, although it is not limited thereto. As seen in Fig. 1, and as discussed in the parent application, the physical layer portion 130 of the ATM destination switch 100 preferably includes a UTOPIA interface 150, a managing RAM 162, a flow controller 166, a microprocessor interface 167, channel interface buffers 170, and a RAM interface 175. The flow controller 166 is coupled to the UTOPIA interface 160, the managing RAM 162, the microprocessor interface 167, the channel interface buffers 170, and the RAM interface 175. The UTOPIA interface generally receives cells of ATM data in a byte-wide format, and passes them to the flow controller 166. Based on the destination of the cell (as discussed in the parent application), and the priority of the cell, the flow controller 166 writes the cell into an appropriate output buffer 170. The output buffer is preferably capable of storing at least two ATM cells so that one cell can be read out of the buffer as another is being read into the buffer without conflict. If buffer space is not available for a particular cell at a particular time, the flow controller 166 forwards the ATM cell via the RAM interface 175 to a desired location in a shared RAM 180 (which may be on or off chip) based on information contained in the managing RAM 162 as discussed in more detail below. When room becomes available in the output buffer 170 for the cell, the flow controller 166 reads the data out of the shared RAM 180, and places it in the buffer 170. In the background, when not receiving data from the UTOPIA interface, and when not reading data from or writing data to the shared RAM 180 or writing data to the buffers, the flow

controller 166 monitors the integrity of the link list structure contained in the managing RAM, as is described in more detail below. In addition, the flow controller 166 can perform various functions in response to microprocessor command received via the microprocessor interface 167.

The managing RAM 162 may serve various functions, including providing information for assisting in the processing of the header of the ATM cell as discussed in the parent application hereto. For purposes of this invention, however, the managing RAM 162, or at least a portion thereof, is preferably provided as a x bit by y word RAM for the purpose of managing $y-2$ link lists which are set up in the shared RAM 180 ($y-2$ equalling the product of w ports times v priorities). Thus, as seen in Fig. 2, a link list information structure for $y-2$ data queues includes: a head pointer, a tail pointer, a block counter, and a queue empty flag for each of the $y-2$ data queues; a free list block pointer, block counter, and queue empty flag for a free list; and a block pointer for the unused blocks of memory. Each head pointer and tail pointer preferably includes a block pointer and a position counter, with the block pointer used for pointing to a block in the memory, and the position counter being used to track pages within a block of memory. Thus, for example, where ATM cells of fifty-three bytes of data are to be stored in the shared memory, and each cell is to be stored on a "page", a block having four contiguous pages may be arranged with the position counter being a two bit counter for referencing the page of a block. The block counter for each queue is used to reference the number of blocks contained within the queue. The queue empty flag when set indicates that the queue is empty, and that the pointers contained within the queue as well as the block count can be ignored.

As suggested above, the head pointer for each link list queue contains the address of the first word of the first memory page of the queue in the shared memory. The tail pointer for each link list queue contains the address of the first word of the last memory page in the queue. Each memory page of the shared

memory is composed of M contiguous memory addresses. Depending on the memory type, each address location can be of size B bits, with common sizes being eight bits (byte), sixteen bits (word), thirty-two bits, or sixty-four bits. In accord with the preferred embodiment of the invention, the address locations are sixteen bits in length with the first M-1 locations in a page containing the stored information. The M'th location of a last page in a block is used to store a next block pointer which is set to the first location of the next block plus an odd parity bit. Where the block is the last block in the queue, the M'th location of the last page in the last block is set to all ones. Where the page is neither the last page of the block, nor the last block in the queue, the M'th location of the page is not utilized. In the preferred embodiment of the invention used with respect to ATM telecommunications data, each page is thirty-two words in length (i.e., $M = 32$), with each word being sixteen bits. Thus, an ATM cell of fifty-three bytes can be stored on a single page with room to spare. It should be appreciated, that in some applications, only the data payload portion of the ATM cell (i.e., forty-eight bytes), and not the overhead portion (five bytes) will be stored in the shared memory. In other applications, such as in switches where routing information is added, cells of more than fifty-three bytes may be stored. Regardless, with a thirty-two word page, system addressing is simplified.

An example of the memory organization of the shared memory is seen in Fig. 3a. In Fig. 3a, two active link list data queues are represented, as well as a free list queue and an Unused block. In particular, memory blocks 512, 124, and 122 are shown linked together for a first queue, memory blocks 511, 125, and 123 are linked together for a second queue, memory blocks 510 - 125 are linked together for the free list queue, and memory blocks 121 - 1 are Unused. It will be appreciated that in the preferred embodiment of the invention, each page contains thirty-two sixteen bit words. Thus, the thirty-second (M'th) word of memory block 512 (seen in more detail in Fig. 3b) contains a pointer (the ten least significant bits) which points to memory block 124, the thirty-second word of memory block 124 contains a

pointer which points to memory block 122, and the thirty-second word of memory block 122 contains all ones, thereby indicating the last word in the queue. Likewise, the thirty-second word of memory block 511 contains a pointer which points to memory block 125, the thirty-second word of memory block 125 contains a pointer which points to memory block 123, and the thirty-second word of memory block 123 contains all ones, thereby indicating the last word of that queue.

The free list of Fig. 3a is seen extending from block 510 to block 126. The unused blocks run from block 121 to block 1.

Turning to Fig. 3c, specifics are seen of the management RAM which would be associated with managing the shared memory in the state of Fig. 3a. In particular, information for link list #1 is seen with a head pointer having a block pointer having a value equal to 512 and a position counter set at "00" to indicate a first page of memory in the block storing data. The tail pointer of the link list #1 information has a block pointer having a value equal to 122 and a position counter set to "11" to indicate that all pages of block 122 are being used. The block counter of the information for link list #1 is set to a value of three, and the queue empty flag is not set (i.e., equals zero). Information for link list #2 is seen with a head pointer having a block pointer having a value equal to 511 and a position counter set at "01" to indicate that the data first occurs at a second page of the block (i.e., the first page already having been read from the block). The tail pointer of the link list #2 information has a block pointer having a value equal to 123 and a position counter set at "10" which indicates that there is no data in the last page of the block. The block counter of the link list #2 information is also set to a value of three, and the queue empty flag is not set. The value of the head and tail pointers and block count for the information of link list #N are not indicated, as the queue empty flag of link list #N is set (equals one), thereby indicating that the pointers and block counter do not store valid data. Likewise, while details of information for other link lists are not shown, the only data of interest would

be that the queue empty flags related to all of those link lists would equal one to indicate that no valid data is being stored with reference to those link lists. The head pointer of the free list information has a block pointer set to a value 510, and a block count of 385. The queue empty flag of the free list is not set, as the free list contains data. Finally, the block pointer relating to the Unused queue is shown set to a value of 121. It is noted that in order to increase performance, the free list head pointer and block counter information is preferably implemented in a series of flip-flops, and is thus readily available for purposes discussed below with reference to Figs. 4a-4d. The queue empty flags are also preferably similarly implemented.

It should be appreciated that by providing the queue empty flags and an Unused block pointer, excessive initialization requirements are eliminated. As suggested above, the queue empty flag indicates that there is no valid data for a link list and that the head and tail pointers for that link list and the block counter of that link list can be ignored. The Unused block pointer is provided to point to the next unused block in shared memory. As memory pages are written or used, the Unused block pointer is decremented until a value of zero is reached. At that point, all cell blocks are included in the link lists (including the free list). As previously mentioned, when a block is read from the shared memory, the available block is added to the free list. When a new block is required for adding to a link list, the block space may be taken from either the free list or from the Unused blocks, and available blocks from the free list may be taken either before or after the Unused blocks are used.

Turning now to Figure 4a, a flow chart of operations of the flow controller 166 of the apparatus 100 of Figure 1 is seen with respect to writing data to the shared memory. It is noted that while the operations are shown in flow chart form, in accord with the preferred embodiment of the invention, the operations are carried out in hardware. When the flow controller 166 determines that it is receiving an ATM cell which cannot be written into a

buffer directly, the flow controller makes a determination at 200 (by checking the management RAM queue empty flag associated with that queue) as to whether the queue which should receive that cell is empty. If the queue is not empty, at 202 the queue status (i.e., the tail pointer and position counter) for that queue is obtained, and at 204 a determination is made as to whether a new block will be needed to be added to the queue (i.e., is the position counter equal to "11"). If a new block is not required, at 206 the cell is written to the shared RAM location indicated by the tail pointer position counter for that queue (stored in management RAM), and at 208 the tail pointer position counter for that queue is updated. At 210, a determination is made as to whether the cell is being written into the last page of a block. If so, at 212 the flow controller writes a word of all ones into the M'th location of the page (in the shared memory).

If it is determined that a new block of shared RAM is required to store the incoming cell because at 200 the queue was empty, at 214, a block is obtained from either the free list or from unused RAM. If the block is obtained from the free list, at 216, the free list information is updated by changing the head pointer of the free list (i.e., setting the head pointer to the value stored in the M'th location of the last page of the obtained block), and by decrementing the block counter associated with the free list. If the block is obtained from the unused RAM, the block pointer for the unused RAM is decremented at 216. Regardless, at 218, the cell is written to the queue, and at 220, the tail pointer and block counter for the queue are both updated in the management RAM (with the block counter being set to the value one), and the queue empty flag is changed.

If it is determined that a new block of shared RAM is required to store the incoming cell because at 204 the tail pointer position counter of the link list indicated that the entire tail block is storing data, at 222, a block is obtained from either the free list or from unused RAM. If the block is obtained from the free list, at 224, the free list is updated by

changing the head pointer of the free list (i.e., setting the head pointer to the value stored in the M'th location of the last page of the obtained block), and by decrementing the block counter associated with the free list. If the free list becomes empty because a block is removed, the queue empty flag of the free list is set. If the block is obtained from the unused RAM, the block pointer for the unused RAM is decremented at 224. Regardless, at 228, the cell is written to the queue, and at 230, the tail pointer and block counter for the queue are both updated in the management RAM.

The details of the flow controller operation with respect to a cell read operation (i.e., where a cell is to be read from a queue because a buffer is available to receive the cell) is seen in Fig. 4b. In particular, when a data buffer becomes available, the flow controller at 250 reads the head pointer and tail pointer in the management RAM for the link list associated with the available data buffer. Then, at 252, the flow controller reads from shared memory the cell at the location in the shared memory indicated by the head pointer, and provides the cell to the data buffer. After the data has been read, the flow controller determines at 254 (based on the head pointer and tail pointer) whether the cell was the last cell in the queue, and at 256 (based on the head pointer position counter) whether the cell was the last cell in a block. If it is neither, then at 258 the queue status is updated (i.e., the head pointer position counter is changed), and another cell read operation is awaited. If at 254 it is determined that the cell is the last cell in the queue, at 260, the head pointer for the free list (obtained from the management RAM) is inserted into the last word of the last page of the freed block. Then at 262, the free list in the management RAM is updated by adding the freed block to the head of the free list; i.e., by updating the free list block pointer and block counter. At 264, the queue empty flag is set for the link list which now has no blocks. If the free list was empty prior to adding the freed block, the free list must be initialized (with appropriate head pointer and block counter) and the queue empty flag changed at 264. In addition, in the case were the free list

was empty prior to adding the freed block, the last word in the freed block in the shared RAM should be set to all ones.

If at 256 it is determined that the cell which has been read out of shared memory is the last in a block, then at 266, the head pointer for the free list as obtained from the management RAM is inserted into the last word of the last page of the freed block. Then, at 268, the queue status for the link list is updated by changing the block pointer and position counter of the head pointer (to the value contained in the last word of the page of memory being read out of the shared memory), and by decrementing the block counter. Again, it is noted that if the free list was empty prior to adding the freed block, the free list must be initialized (with appropriate head pointer and block counter) and the queue empty flag changed, and the last word in the freed block in the shared RAM should be set to all ones. It is also noted, that upon obtaining the pointer in the M'th location of the last page of the block, according to the preferred embodiment of the invention, at 270, a parity check is done on the pointer. At 272, the calculated parity value is compared to the parity bit stored along with the pointer. Based on the comparison, at 274, a parity error condition can be declared, and sent as an interrupt message via the microprocessor interface port 167 (Fig. 1) to the microprocessor (not shown). Preferably, when a parity error is found, the microprocessor treats the situation as a catastrophic error and reinitializes the management and data RAMs.

Figure 4c sets out the operation with respect to the queue clear microprocessor command (received via the microprocessor interface 167). When the queue clear operation is presented, at 270 the queue status for the link list is updated by setting the queue empty flag, and at 272, the blocks in the queue are added to the head of the free list which is updated in a manner discussed above (Fig. 4b) with reference to the cell read operation.

The link list monitoring operation seen in Fig. 4d is the hardware agent which monitors the integrity of the link list structure whenever the cell write, cell read, and queue clear operations are not running. The link list monitoring operation preferably monitors four different error conditions. In particular, at 280, the counts of all of the link list block counters (including the free list) where the queue empty flag for those link lists are not set are summed together with the unused blocks and compared the total number of memory blocks in the common RAM. If the sum does not equal the total number of memory blocks in the common RAM, at 281, an error condition is declared by triggering a microprocessor interrupt bit. At 282, the head and tail block pointers of each link list are compared. If at 284 the head and tail block pointers are determined to be equal, at 286 the block counter is checked, and if not equal to one, at 287 an error condition is declared. If the head and tail block pointers are not equal when compared at 284, at 288 the block counter is checked, and if the block count is equal to one, at 289 an error condition is declared. At 290, the block counter for each link list whose queue empty flag is not set is checked; and if the block counter equals zero, at 291 an error condition is declared.

According to the preferred embodiment of the invention, the write, read, clear, and monitoring operations of the flow controller are carried out in hardware which may be generated by using HDL code to synthesize hardware gates via use a VHDL compiler. Figures 5a-5d are state machines diagrams corresponding to the HDL code, including a write state machine (Fig. 5a), a read state machine (Fig. 5b), a clear state machine (Fig. 5c), and a monitoring state machine (Fig. 5d). The gates created using the code may be standard cell technology or gate array technology.

It should be appreciated that the invention is not intended to be limited to a strictly hardware implementation, but is also intended to apply to memory management utilizing a microprocessor with associated firmware (e.g., a ROM).

There have been described and illustrated herein an apparatus and method for management of shared memory. While particular embodiments of the invention have been described, it is not intended that the invention be limited thereto, as it is intended that the invention be as broad in scope as the art will allow and that the specification be read likewise. Thus, while the invention has been described with reference to VLSI implemented ATM equipment, it will be appreciated that the invention has broader applicability. Also, while specific details of RAM sizes, etc. have been disclosed, it will be appreciated that the details could be varied without deviating from the scope of the invention. For example, while a management of RAM of size x bits by y words has been described for managing $y-2$ link lists of data, it will be appreciated that the management RAM could assume different sizes. Thus, for example, instead of using a separate word for the unused block pointer, the unused block pointer could be located in the "tail pointer" location of the free list (which itself does not use a tail pointer), thereby providing a management RAM of x bits by y words for managing $y-1$ link lists of data. In addition, rather than providing the information related to the link lists with the head pointer, tail pointer, block counter, and queue empty flag in that order, the variables of the link list could be reordered. Similarly, instead of providing a shared memory having pages of thirty-two words in depth, each word being sixteen bits in length, it will be appreciated that memories of different lengths and depths could be utilized. Also, rather than locating the pointer to the next block in the last word of the last page of a previous block, it will be appreciated that the pointer could be located in a different location. Further yet, while specific flow charts have been disclosed with respect to various operations, it will be appreciated that various aspects of the operations can be conducted in different orders. In addition, while particular code has been disclosed for generating gate arrays which conduct the operations in hardware, it should be appreciated by those skilled in the art that other code can be utilized to generate hardware, and that hardware and/or firmware can be generated in

different manners. Furthermore, while the invention was described with respect to separate RAMs for the management RAM and the shared data RAM, it will be appreciated that both memories may be part of a larger single memory means. It will therefore be appreciated by those skilled in the art that yet other modifications could be made to the provided invention without deviating from its spirit and scope as so claimed.

Claims:

1. Apparatus for managing the storage of data in a memory, comprising:

a) a shared memory means having a plurality of data storage locations;

b) control means for receiving said data and forwarding said data to desired of said plurality of data storage locations in said shared memory means, wherein said data is stored in said plurality of data storage locations in the form of a plurality of link lists, each link list having a head;

c) management memory means for storing information regarding each of said plurality of link lists, said information including a head pointer and a queue empty flag for each link list, said head pointer for each particular respective link list pointing to a location of a respective said head of that particular link list, and said queue empty flag for a link list indicating that that link list has no valid data contained therein.

2. An apparatus according to claim 1, wherein:

said control means reads data from said shared memory means, at least a plurality of said data storage locations are in the form of a free link list, said free link list relating to data storage locations from which data has been read by said control means, and

said management memory means includes a pointer and a queue empty flag for said free link list.

3. An apparatus for managing the storage of data in a memory, comprising:

a) a shared memory means having a plurality of data storage locations;

b) control means for receiving said data and forwarding said data to desired of said plurality of data storage locations in said shared memory means, and for reading data from said shared memory means, wherein said data is stored in said plurality of data storage locations in the form of a plurality of link lists, each link list having a head;

c) management memory means for storing information regarding each of said plurality of link lists, said information including a head pointer for each link list queue, said head pointer for each particular respective link list pointing to a location of a respective said head of that particular link list,

wherein upon initialization, at least a plurality of said data storage locations of said shared memory means are unused, and after utilization, at least a plurality of said data storage locations are in the form of a free link list, said free link list relating to data storage locations from which data has been read by said control means, and

wherein said management memory means includes a pointer to at least one of said unused data storage locations, and said management memory means includes a pointer for said free link list.

4. An apparatus according to any preceding claim, wherein:

at least upon initialization, at least a plurality of said data storage locations of said shared memory means are unused, and

said management memory means includes a pointer to at least one of said unused data storage locations.

5. An apparatus according to any previous claim, wherein:

said shared memory means is arranged in a plurality of blocks with each block having a plurality of said data storage locations, and

said information stored in said management memory means regarding each of said plurality of link list queues includes a block counter for each of said plurality of link list queues, each block counter counting the number of blocks contained in that link list queue.

6. An apparatus according to claim 5, wherein:

each of said plurality of blocks is arranged as a plurality of contiguous pages with each page having a plurality of said data storage locations, and

each said head pointer comprises a block pointer which points to a block and a page counter which points to a page in said block.

7. An apparatus according to claim 5, wherein:

each block storing data includes at least one location containing one of (i) a pointer to a next block in the link list, and (ii) an indicator which indicates that the block is the last block in the link list.

8. An apparatus according to claim 7, wherein:

said pointer to a next block in the link list includes a parity bit for said pointer.

9. An apparatus according to claim 6, wherein:

each block storing data includes at least one location in a last page of that block containing one of (i) a pointer to a next block in the link list, and (ii) an indicator which indicates that the block is the last block in the link list.

10. An apparatus according to any previous claim, wherein:

said information includes a tail pointer for each link list containing said data.

11. An apparatus according to claim 6, wherein:

said information includes a tail pointer for each link list containing said data,

each of said plurality of blocks is arranged as a plurality of contiguous pages with each page having a plurality of said data storage locations,

each said head pointer comprises a first block pointer which points to a block and a page counter which points to a page in said block, and

each said tail pointer comprises a second block pointer which points to a tail block and a page counter which points to a page in said tail block.

12. An apparatus according to claim 6, wherein:

said data comprises ATM data received in cell format, and each said page includes enough of said data storage locations to store all of the data contained in an ATM cell.

13. An apparatus according to claim 12, wherein:

each page includes thirty-two sixteen bit word locations.

14. An apparatus according to claim 5, wherein:

said control means reads data from said shared memory means, at least a plurality of said data storage locations are in the form of a free link list, said free link list relating to data storage locations from which data has been read by said control means, and

said management memory means includes a pointer, a block counter, and a queue empty flag for said free link list,

at least a plurality of said data storage locations of said shared memory means are unused, and

said management memory means includes a pointer to said at least one of said unused data storage locations, and

said control means includes means for comparing a sum of counts of said block counters of each link list containing data, said free link list, and said unused pointer to the number of blocks in said shared memory means.

15. An apparatus according to claim 14, wherein:

said control means further comprises means for generating an error signal is said sum of counts does not equal said number of blocks in said shared memory means.

16. An apparatus according to claim 10, wherein:

said control means includes means for comparing, for each link list containing data, said tail pointer to said head pointer.

17. An apparatus according to claim 16, wherein:

said control means further comprises means for generating an error signal if said tail pointer and said head pointer for a link list containing data point to an identical block, and said block counter for said link list does not equal one.

18. An apparatus according to claim 16, wherein:

said control means further comprises means for generating an error signal if said tail pointer and said head pointer for a link list containing data point to different blocks, and said block counter for said link list equals one.

19. An apparatus according to claim 5, wherein:

said control means further comprises means for checking the count of each block counter of a link list where the queue empty flag is not set, and for generating an error signal if the count is zero and the queue empty flag is not set.

20. An apparatus according to any preceding claim, wherein:

said control means and said management memory means are contained on a single integrated circuit.

21. An apparatus according to claim 5, wherein:

said management memory means includes said pointer, a block counter, and a queue empty flag for said free link list, and

said control means includes means for comparing a sum of counts of said block counters of each link list containing data, said free link list, and said unused pointer to the number of blocks in said shared memory means, and means for generating an error signal if said sum of counts does not equal said number of blocks in said shared memory means.

22. An apparatus according to claim 10, wherein:

said control means includes means for comparing, for each link list containing data, said tail pointer to said head pointer, and means for generating an error signal if either

(i) said tail pointer and said head pointer for a link list containing data point to an identical block, and said block counter for said link list does not equal one, or

(ii) said tail pointer and said head pointer for a link list containing data point to different blocks, and said block counter for said link list equals one.

23. A method of managing the storage of data utilizing a controller, a shared memory having a plurality of data storage locations, and a management memory, said method comprising:

a) using said controller to forward received data to desired of the plurality of data storage locations in the shared memory, wherein the data is stored in the plurality of data storage locations in the form of a plurality of link lists, each link list having a head; and

b) storing information regarding each of the plurality of link lists in the management memory, said information including a head pointer and a queue empty flag for each link list, said head pointer for each particular respective link list pointing to a location of a respective said head of that particular link list, and said queue empty flag for a link list indicating that that link list has no valid data contained therein.

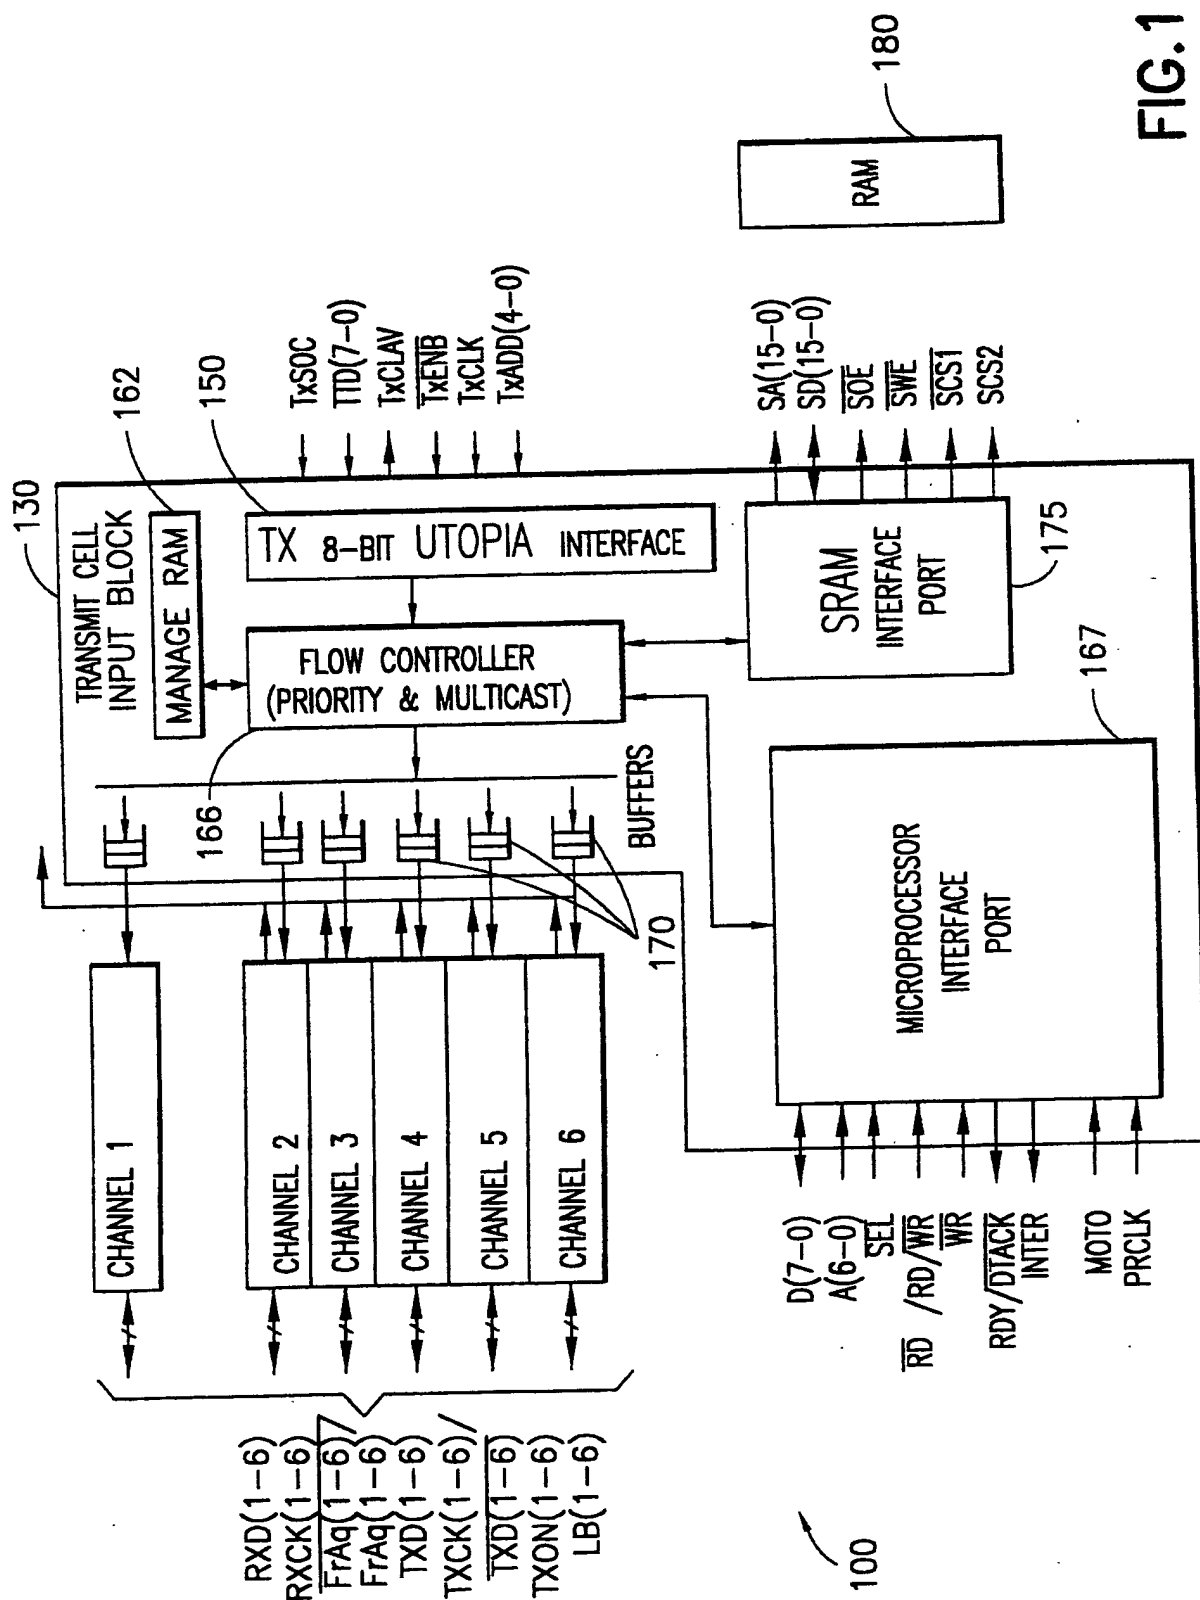


FIG. 1

2/11




FIG.2

3/11




FIG.3a

LAST PAGE IN MEMORY
BLOCK 512

+1	DATA
	⋮
+31	DATA
+32	00000P0001111100

FIG.3b

4/11




FIG.3c

5/11

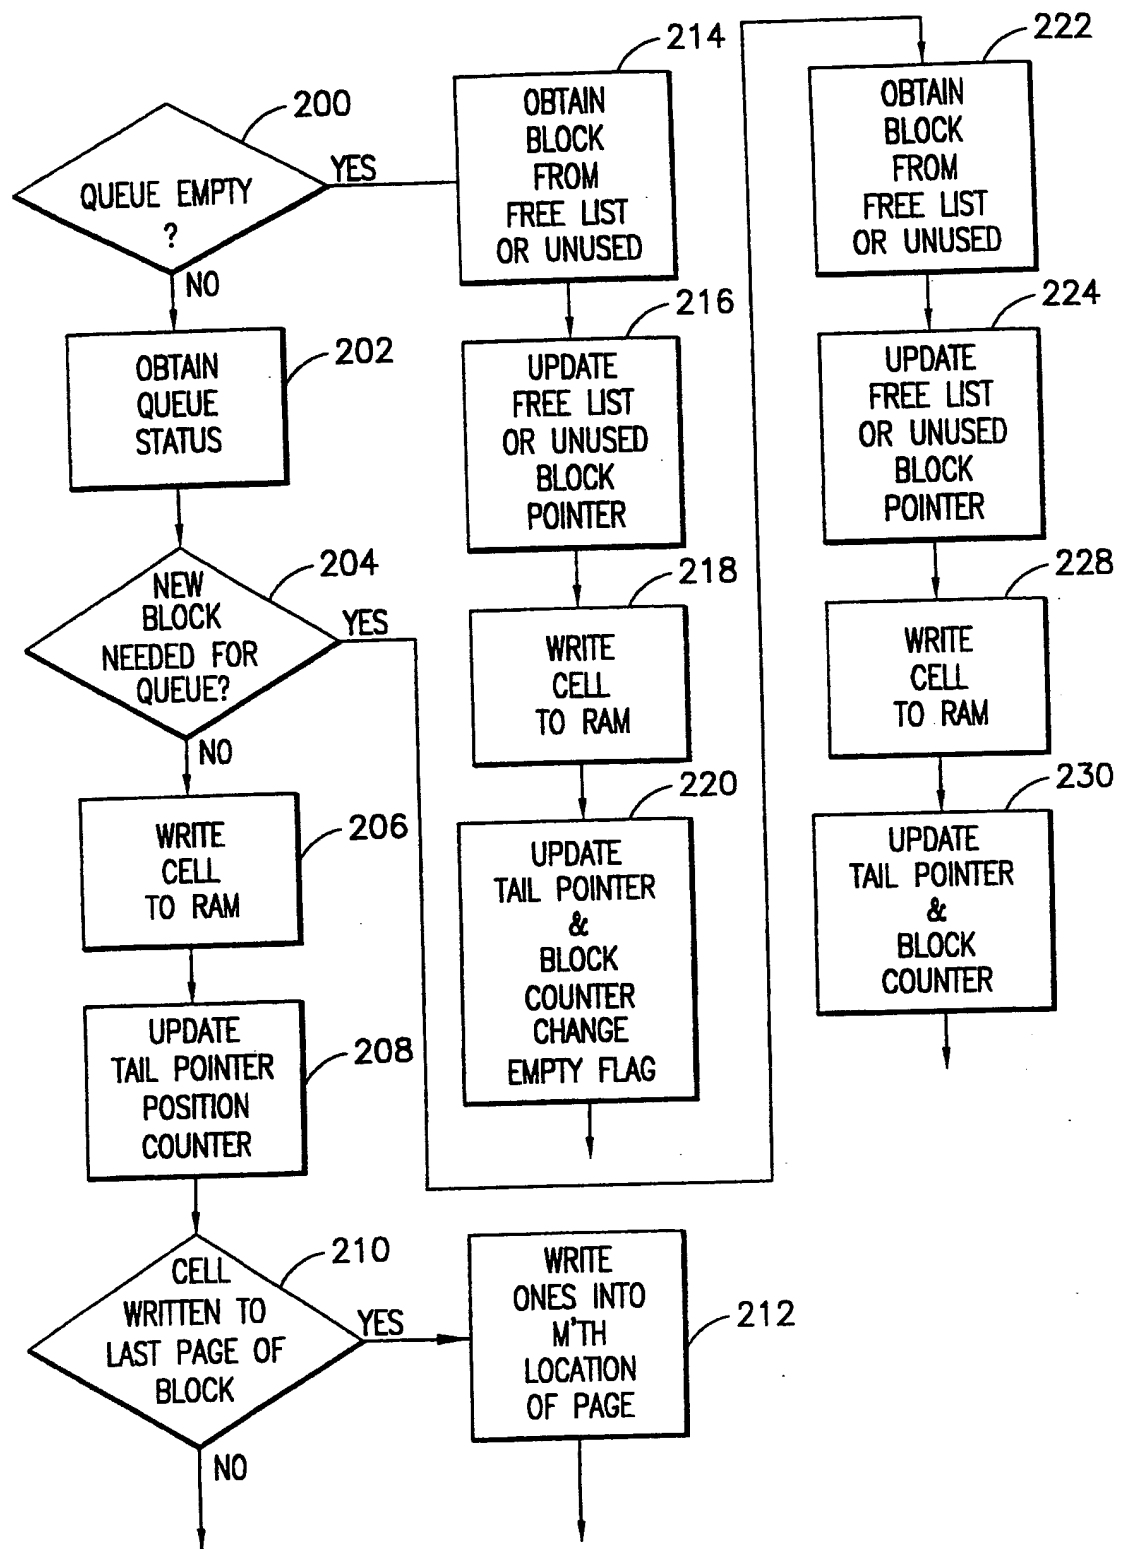


FIG. 4a

SUBSTITUTE SHEET (RULE 26)

6/11

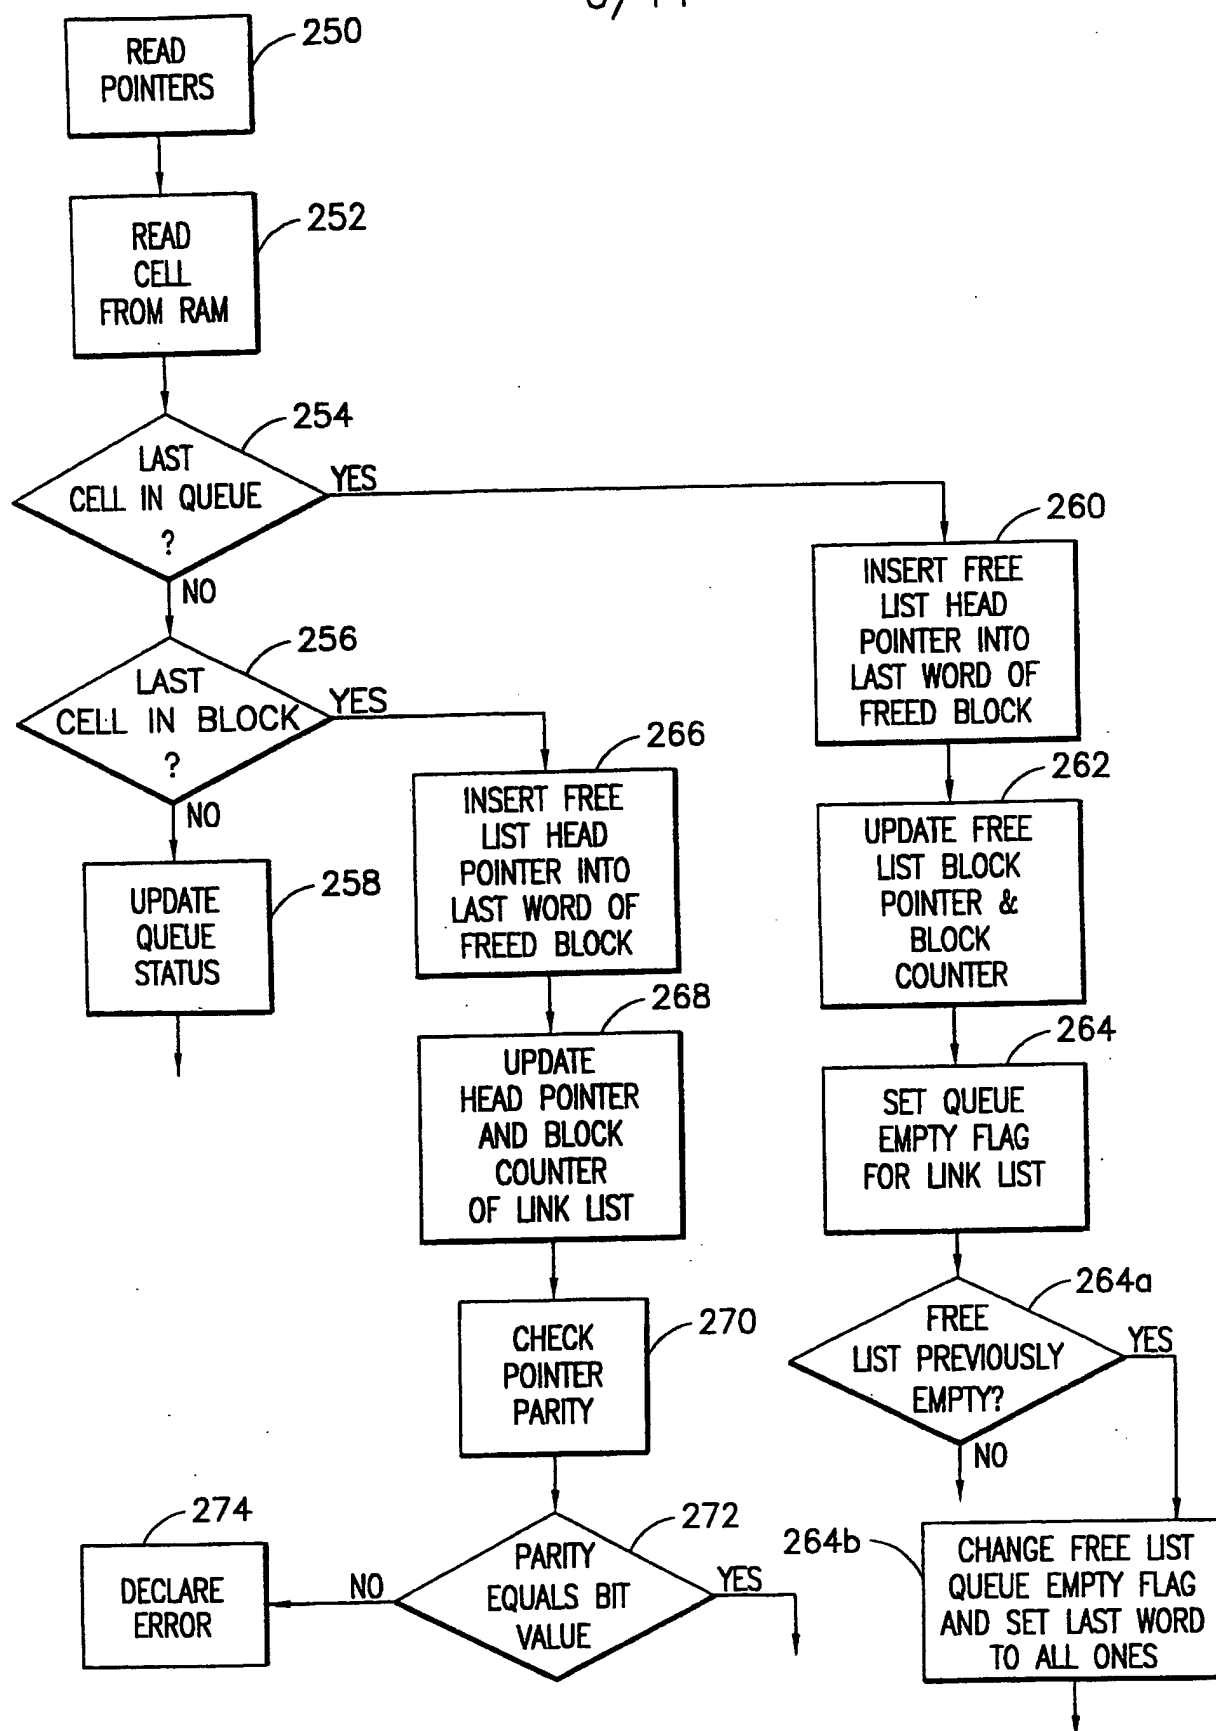


FIG. 4b

SUBSTITUTE SHEET (RULE 26)



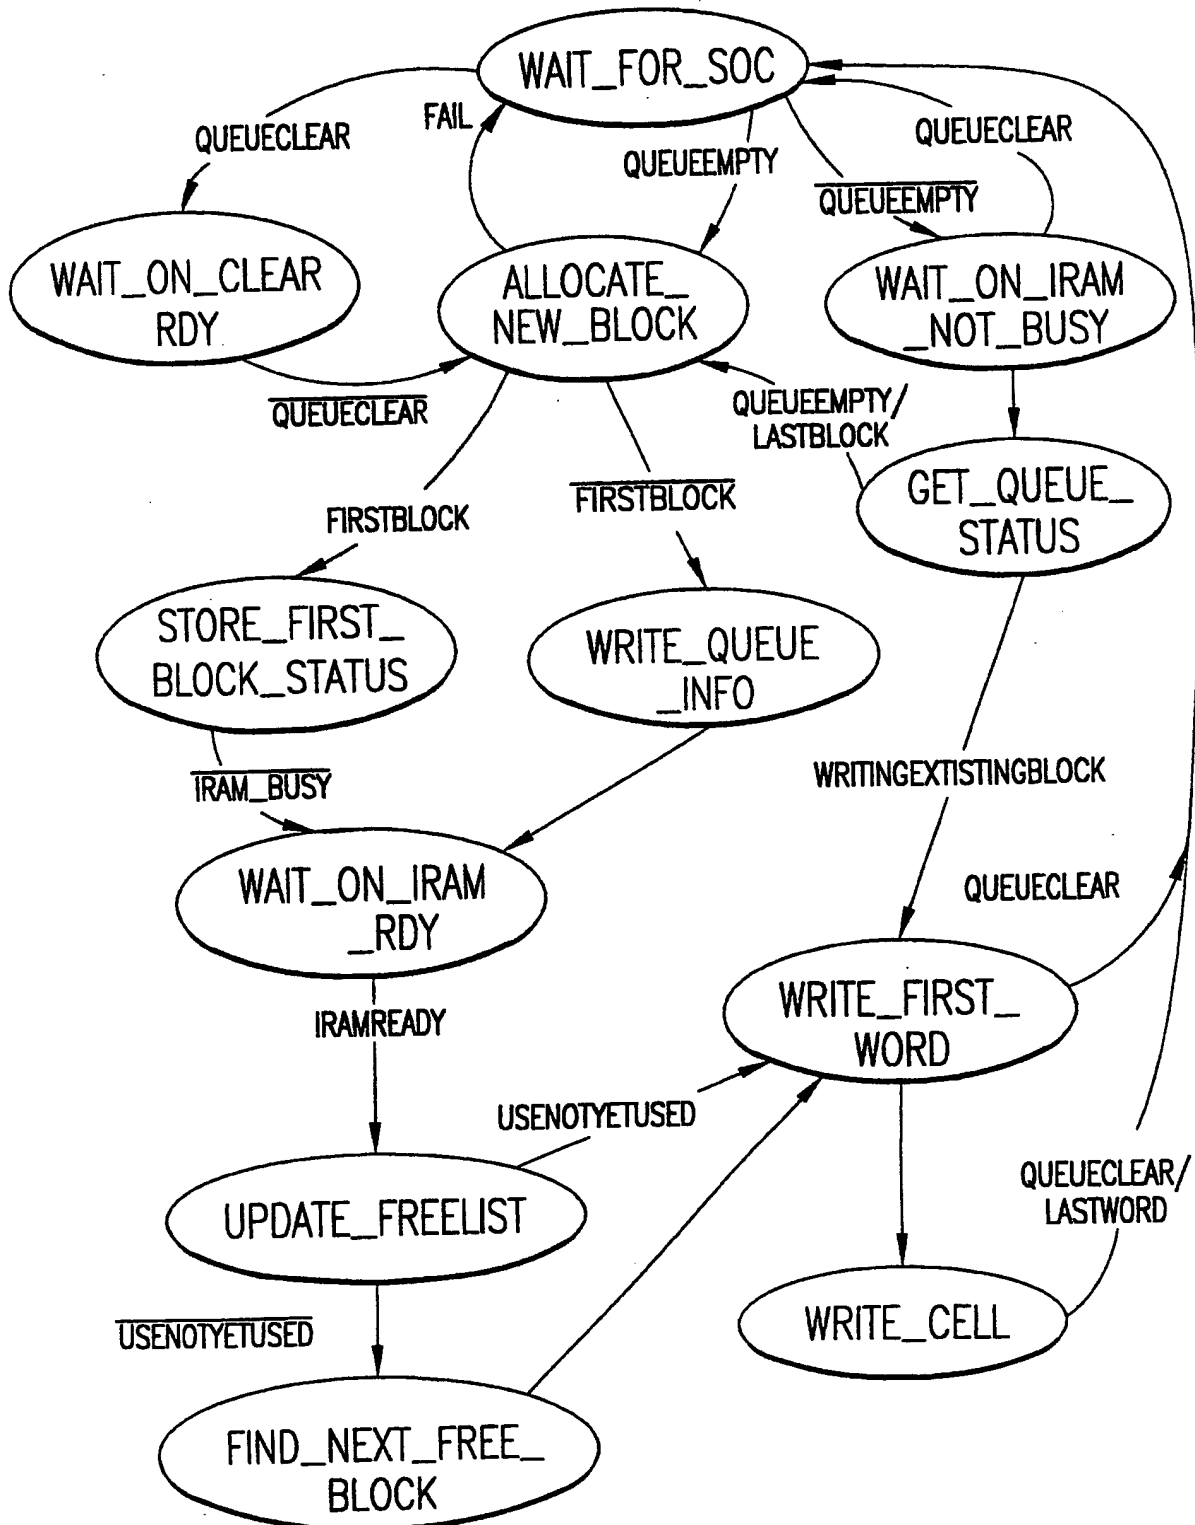


FIG.4d

8/11

WRITE STATE MACHINE**FIG.5a**

9/11

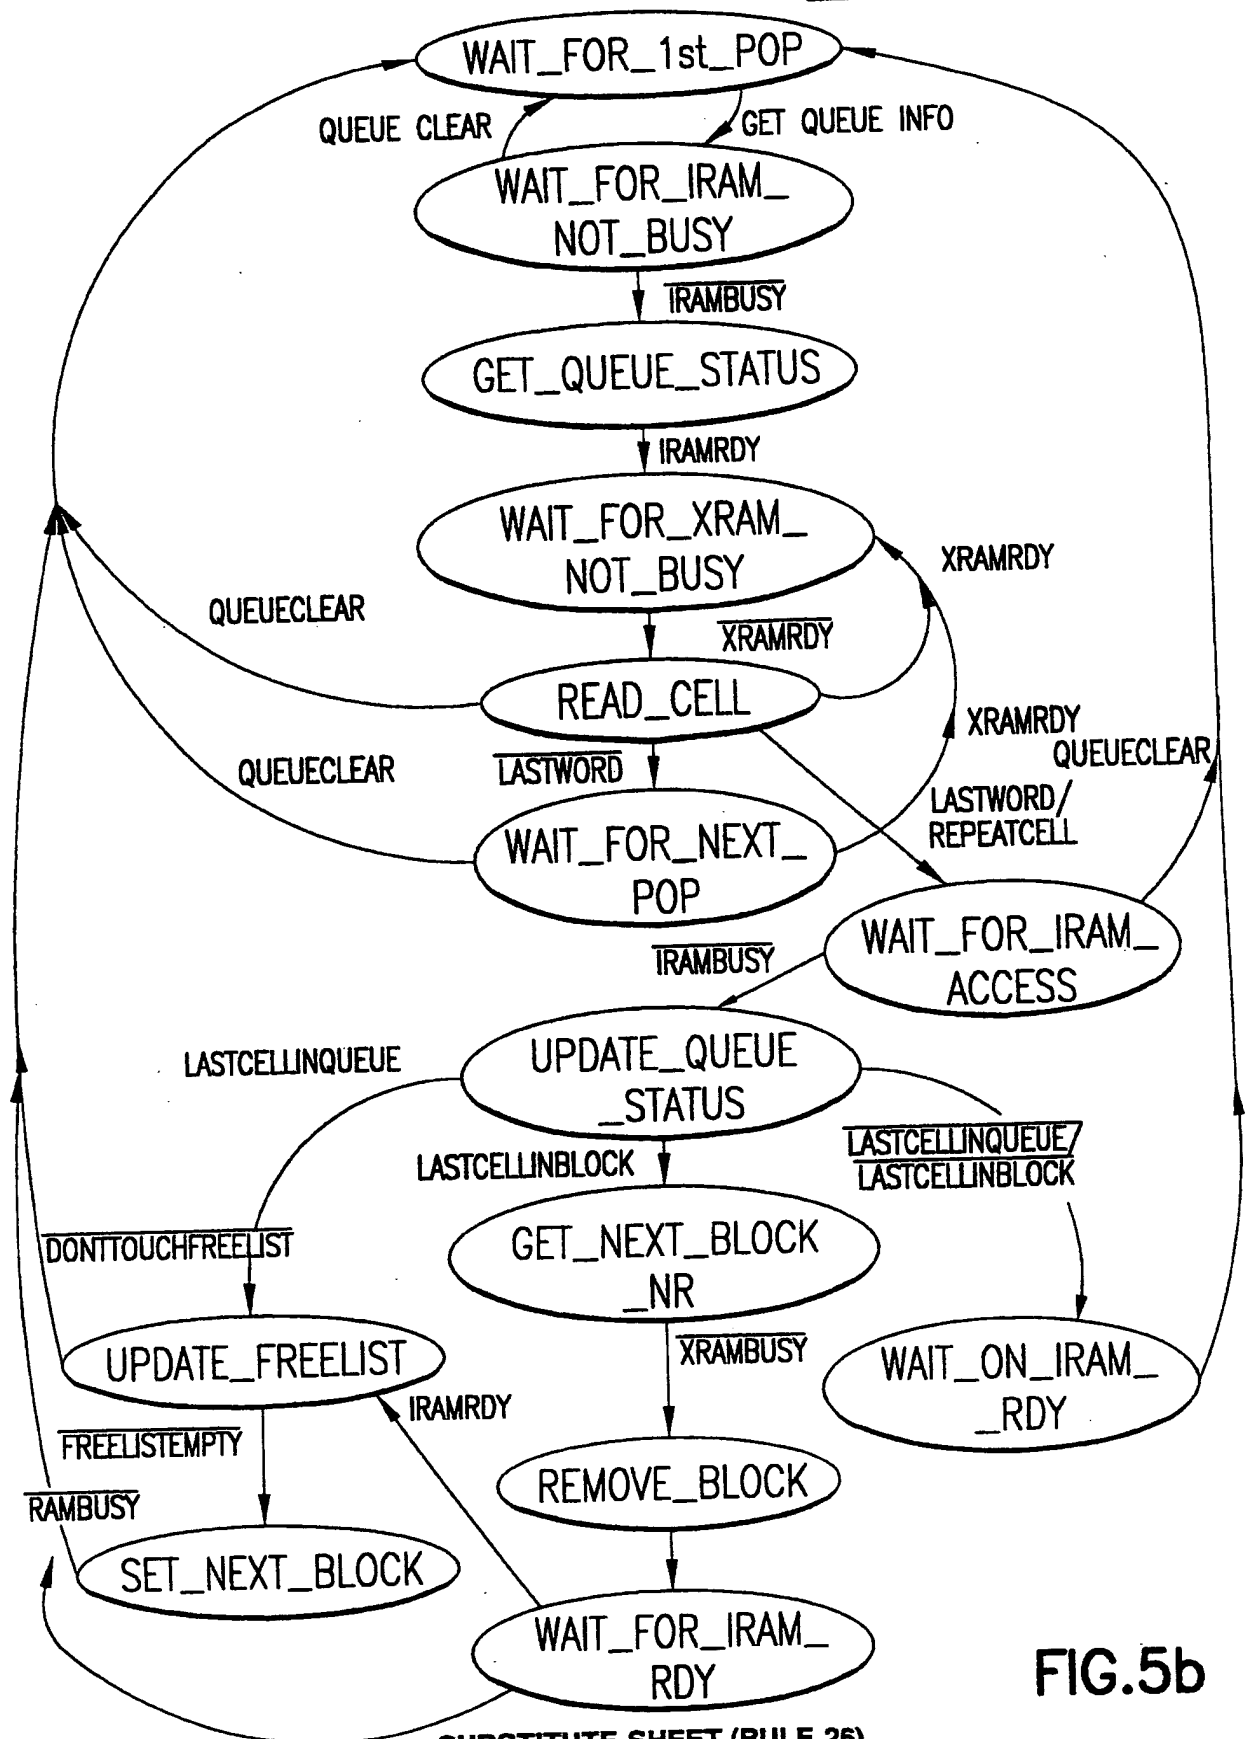
READ STATE MACHINE

FIG.5b

10/11

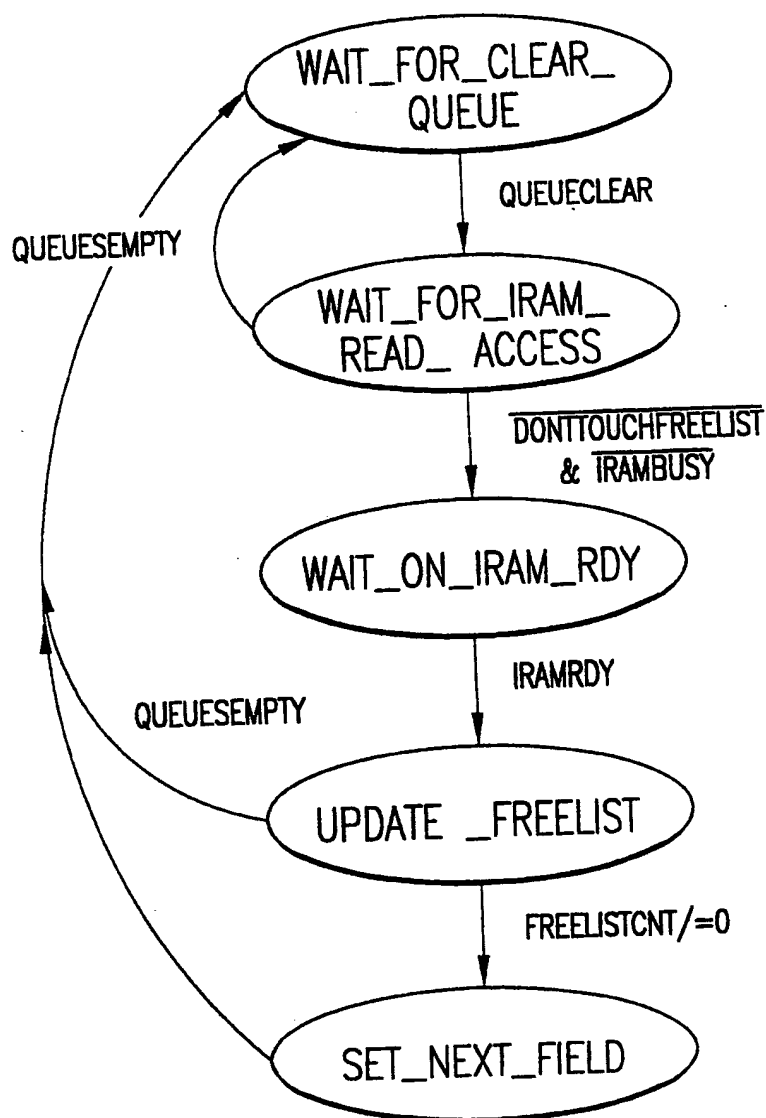
CLEAR STATE MACHINE

FIG.5c

11/11

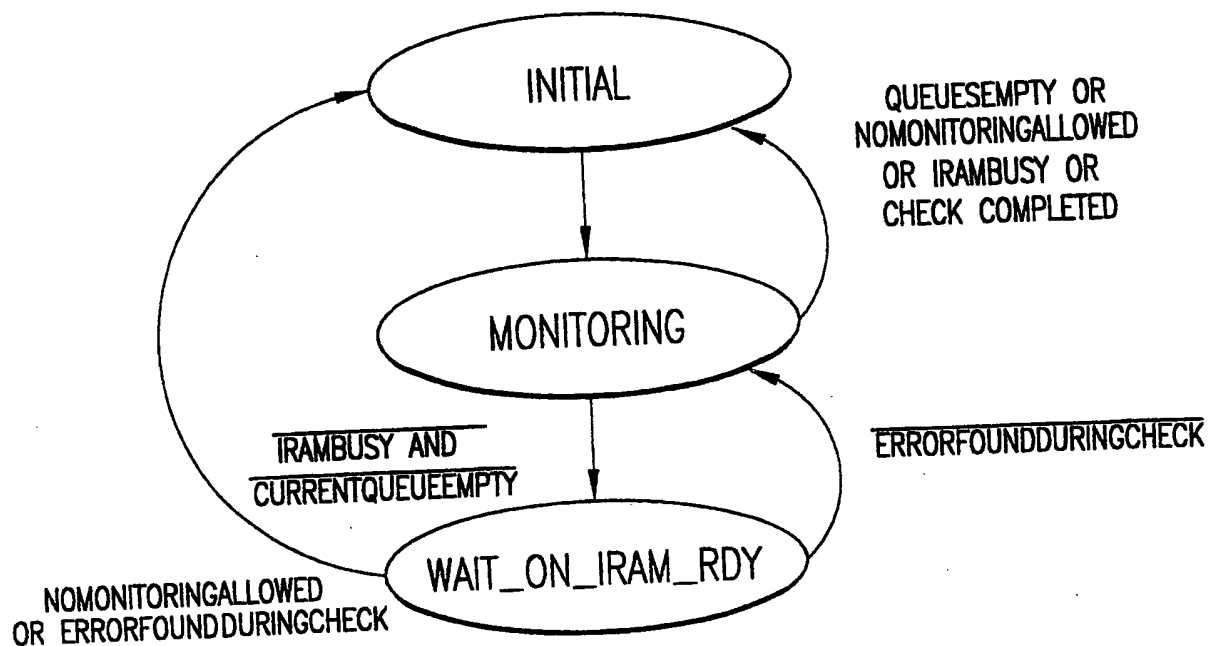
MONITORING STATE MACHINE

FIG.5d

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/02131

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 12/00

US CL :711/153

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 711/153, 711/207, 370/232, 370/398

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, MAYA

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X — Y	US 5,390,175 A (HILLER ET AL.) 14 February 1995, col. 45, line 10; Fig. 24; col. 37, line 20; col. 37, line 25; col. 21, line 35; col. 22, line 37; col 20, line 63; col. 54, line 34; col. 55, line 53; col. 35, line 56; col. 21, line 60; col. 37, line 19; col. 35, line 56	1-5, 7, 8, 10, 14-19, 20, 21, 22, 23 6, 9, 11-13
Y	US 5,123,101 A (SINDHU) 16 June 1992, col. 21, line 68.	6, 9, 11-13
A, P	US 5,654,962 A (ROSTOKER ET AL) 05 August 1997	1-23

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Z* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

30 MARCH 1998

Date of mailing of the international search report

0 4 AUG 1998

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

DAVID LANGJAHR

Telephone No. (703) 305-4034

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☒ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.